

5 SUMMARY

In another embodiment, in case of a link port failure, the restoration is localized to the node where link port failure occurred. If a link port failure occurs between two adjacent nodes then localization is effected through restoration by the two adjacent nodes without the resource utilization of any other node or the network

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a zoned network consisting of four zones and a backbone.

- 3 -

Fig. 3A is a flow diagram of steps performed in response to failure of a port on a link between two adjacent neighbor nodes.

Fig. 3B illustrates an example of a Restore Span packet.

Fig. 3C illustrates an example of a Restore Span Response packet.

5 Fig. 3D illustrates an example of a *Restore_I* packet.

Fig. 4A is a flow diagram of general steps performed by a source or proxy node in response to the failure of a link.

Fig. 4B illustrates an example of a Restore Path Request (RPR) packet.

10 Fig. 5 illustrates the actions performed by tandem nodes in processing received RPR tests.

Fig. 6 illustrates the process performed at the target node once the RPR finally reaches that node.

Fig. 7 Is a flow diagrams illustrating the processes performed by originating nodes that receive negative RPR responses.

15 Fig. 8 Is a flow diagrams illustrating the processes performed by originating nodes that receive negative RPR responses.

Fig. 9 is a block diagram illustrating network environment in a commercial transaction processing system.

20 Fig. 10 is a block diagram of a host computer system suitable for implementing the present invention.

Fig. 11 illustrates the interconnection of host computer system to client systems.

The use of the same reference symbols in different drawings indicates similar or identical items.

25

DETAILED DESCRIPTION OF THE INVENTION

30 The following is intended to provide a detailed description of an example of the invention and should not be taken to be limiting of the invention itself. Rather, any number of variations may fall within the scope of the invention which is defined in the claims following the description.

Network architecture

To limit the size of the topology database used by some protocols, and to limit the scope of broadcast packets (e.g., packets used for restoration), the nodes of a network can be divided into logical groups, referred to herein as “zones.” The use of zones provides several benefits. Networks employing zones can be implemented in several different ways, some of which can be implemented concurrently.

Fig. 1 illustrates an exemplary network that has been organized into a backbone, zone 100, and four configured zones, zones 101-104, which are numbered 0-4 under the protocol, respectively. The exemplary network employs a type 0 node ID, as there are relatively few zones (4). The solid circles in each zone represent network nodes, while the numbers within the circles represent node addresses, and include network nodes 111-117, 121-126, 131-136, and 141-147. The dashed circles represent network zones. The network depicted in Fig. 1 has four configured zones (zones 101-104 (addressed as zones 1-4) and one backbone (zone 100 (addressed as zone 0)). Nodes with node IDs 1.3, 1.7, 2.2, 2.4, 3.4, 3.5, 4.1, and 4.2 (network nodes 113, 117, 122, 124, 134, 135, 141, and 142, respectively) are boundary nodes because they connect to more than one zone. All other nodes are interior nodes because their links attach only to nodes within the same zone. Zone 100 consists of 4 nodes, zones 101-104, with node IDs of 0.1, 0.2, 0.3, and 0.4, respectively.

Another example of the use of zone boundaries is in the provisioning and restoration of circuits within the network. Zone boundaries can be used to limit the flow of information generated by certain nodes during such provisioning and restoration. For example, a node can act as a proxy node for the source or destination node. In the event of a failure in the network affecting the circuit between the two nodes, the proxy node can perform source or destination node related functions. In the case of a failure, a node at the boundary of the zone, in which the failure has occurred, acts as a proxy for the source (or destination) node in the other zone. This avoids the need to involve the portion of the network circuit lying outside of the zone experiencing a failure, which would be expected to remain unchanged.

Another example of localized restoration is the use of the nodes where the physical port failure between two adjacent nodes has occurred. Upon discovering port

5

Once a network topology has been defined, the user can configure one or more end-to-end communication connections that can span multiple nodes and zones, an operation that is referred to herein as provisioning. Each set of physical connections that are provisioned creates an end-to-end communication connection between the two end nodes that supports a virtual point-to-point link (referred to herein as a virtual path or VP). Each link comprises a plurality of physical port between two adjacent nodes. The resulting VP has an associated capacity and an operational state, among other attributes.

20 With each VP is associated the notion of a Physical Instance. The first time that a VP is set up, the VP is assigned a Physical Instance of 1. Whenever a fully set up path for a VP has to be changed because of some failure, the Physical Instance is incremented by 1. The source node of VP maintains a field referred to as Physical Instance Identifier. This field is preferably included in restoration related and other
25 configuration packets. Only the source node should be provided the ability to update this parameter of a VP. The first path of any VP that is successfully provisioned (as seen by the Source node) will typically have a Physical Instance identifier of 1. All future Path related messages (Restore_I, Restore_Path and like) will preferably have the correct value of this identifier. If a new path is selected for this VP, the source
30 node of the VP, which originates Restore Path Request packets, will increment this identifier by 1. As will be discussed subsequently, several Physical Instances of the same VP may temporarily exist in the network at any given time due to the distributed

5

Failure detection, propagation, and restoration

10

15

25

maximum of 125 microseconds to forward the AIS upstream, which quickly propagates failures toward the source (or proxy) node.

Once a node detects a failure on one of its links, either through a local LOS defect or a received AIS indication, the node scans its VP table looking for entries for VPs that include the failed link and analyzes the mode of failure.

Failure Modes:

There can be many virtual path-related failure modes in a network depending upon the location of link failure in that network. Virtual path failure modes can be characterized as follows:

Neighbor or Link mode: In this mode, the failure occurs at any physical port of a link between two adjacent nodes. The scope of the notification and restoration messages is preferably limited to the two adjacent nodes that discover the failure. Restoration, in this case, is relatively fast because the restoration does not involve the resource of any other node in the network (including proxy or source node). The inter-nodal physical path traversed by the VP also does not change.

Path mode: In this mode, path failures are much more broader than a physical port breakdown. This may include a failure of entire optical link or a node. The source and destination of the affected VPs are notified and potentially the entire network is involved in the restoration. This mode may involve more messaging and longer restoration time than the Link mode. However, it depends on the CoS of the VP that how fast the path will be restored. Restoration here will most likely change the path traversed by the VP.

Failure Restoration:

Fig. 2 is a flow diagram of broad general steps performed by a node for restoration in response to the failure of a virtual path. As noted, the failure of a link, results in a LOS condition at the nodes connected to the link and generates the appropriate AIS and RDI indications. If an AIS or RDI was received from a node, a failure has been detected (step 200). In that case, each affected node performs several actions in order to maintain accurate status information with regard to the VPs

5

10

20

25

5

10

15

20

25

Fig. 3D illustrates an example of a *Restore_I* packet. The *Restore_I* packet (360) contains the Identification of failed VP (365), Physical Instance Identification of the VP (370) that the node originating the *Restore_I* packet has and the attempt count of the current physical instance of the Virtual Path (375). The *Restore_I* packet also

contains the direction of the failure (380). The direction field indicates whether the failure occurred on the UPSTREAM (towards the Source of the VP) or DOWNSTREAM (towards Destination) side of the node. The failed path field (385) indicates whether the path is the Primary or Secondary path for CoS-3. For all other
 5 types of VPs, this field indicates that failed path is the primary path for failed VP.

Initiating Restore_I Request: The node that discovers the failure transmits *Restore_I* packets both upstream and downstream. The source of the VP only transmits them downstream. After transmitting packet, the source node immediately starts restoration process. The destination node sends packets upstream and after
 10 sending these packets, the node starts a timer for acknowledgement and retransmission.

The intermediate or tandem node attempt to send packets in both directions even in the case that there is no transmission path in one of the directions. *Restore_I* packet transmissions are reliable and preferably are acknowledged by the receiving
 15 node (the neighbor of the initiating node). A timer is started for each *Restore_I* request that is sent. The request is resent periodically until a response is received or a threshold for the number of unacknowledged packets sent is reached. In either case, if an acknowledgment is received or the threshold is reached, then all resources associated with the VP are released. In the case of multiple port failure in the same
 20 VP, even if *Restore_I* messages are sent out regarding the failure of the first port in a VP, upon subsequent failures, similar procedure is carried out. Figure 3C illustrates a format of the *Restore_I* message.

Receiving Restore_I Request:

Upon receiving a *Restore_I* packet, certain tests are performed. A node should
 25 examine the VP ID and return a NAK (with a reason code of NULL_VP) if this VP doesn't exist at the node. For VP with a lower CoS, the state of the VP and Path is changed to DOWN. When source node of a lower CoS VP receives a *Restore_I* packet, then the source node will generate an alarm. For VPs with a CoS requiring Dynamic Mesh Restoration, if the receiving node is a tandem node, then the current
 30 Physical Instance (PI) of the VP at the node is compared to the PI in the *Restore_I*

09:45:12 PM 11/15/2000

packet. If the stored PI at the node is higher than the one in the *Restore_I* packet, then a NAK (STALE_MSG) is returned (to the sending node). Otherwise, the *Restore_I* packet is accepted for processing. The state of the VP is changed to DOWN, if the state is not already set in that state. The *Restore_I* packet is then forwarded in the same direction as the incoming packet. A timer is initiated for retransmission and after certain predefined threshold times retransmissions, all the resources of VP are released.

If the receiving node is the Source node of the VP, the source node starts restoration and the state of path is set to RESTORING. If the receiving node is the destination node then all the resources associated with the VP are released and the node waits for messages from the Source to arrive. In all the cases (except when a NAK has to be returned), an ACK is returned to the node sending the packet.

Receiving a Restore_I Response:

If the receiving node has already responded to *Restore_I* packet then responses to subsequent *Restore_I* packets are ignored if the Physical Instance of the VP in the response is lower than the node's current Physical Instance for that VP. For valid responses, the appropriate timer and retransmissions are stopped. If receiving node is tandem node and the response is positive and responses to all the packets sent out for that VP have been received, then the resources allocated to the VP are released.

For negative responses, the reason code is examined and if the number of retries is below a given threshold, the sending node resends the *Restore_I* packet. If this threshold has been exceeded, resources associated with the VP are released. Similarly, if the reason code is STALE_MSG (possibly because the responding node has already participated in the set up of a new restoration path for the VP and thus has a higher Physical Instance now), resources for the VP are released without any retransmissions.

Initiating Path Restoration:

Processing by Source / Proxy node:

Fig. 4A is a flow diagram of general steps performed by a source or proxy node in response to the failure of a link. As noted, the failure of a link results in a LOS condition at the nodes connected to the failed link and generates appropriate AIS and RDI indications. When an AIS or RDI message is received from a node, the message indicates detection of a failure (step 400). In that case, each affected node performs several actions in order to maintain accurate status information with regard to the VPs currently supported. The first action taken in such a case is that the node scans its VP table, looking for entries that have the failed link in their path (steps 405). Next, the node checks if the VP uses the failed link (step 410). If the VP does not use the failed link, the node goes to the next VP in the table (step 415). If the selected VP uses the failed link, the node changes VP's state to RESTORING (step 420) and releases all link bandwidth allocated to that VP (step 425). The node places VP on the list of VPs to be restored (step 430).

Next, the source or proxy node sends a Restore Path Request packet (RPR) to all eligible neighbors (i.e., a node adjacent to the given node) in order to restore the given VP. Neighbor eligibility is determined by the state of the neighbor, available link bandwidth, current zone topology, location of the target node, and other parameters.

Fig. 4B illustrates an example of a Restore_Path Request (RPR) packet. The Restore Path packet is sent by the source nodes (or proxy boundary nodes), to obtain an end-to-end path for a VP. The packet is usually sent during failure recovery procedures but can also be used for provisioning new VPs. The RPR packet (440) includes a virtual path identifier field (445) and a reserved flag field (450). The requested bandwidth for failed virtual path is indicated in bandwidth field (455). The Path Length field (460) indicates the number of links in the VP. This field determines how many Link IDs appear at the end of the packet. A HOP_COUNT field (465) is

5

10

15

20

30

forwarded requests. Table 1 lists an example of the fields that are preferably present in an RPRE. Other path relevant information can also be added to the structure.

Field	Usage
Origin Node	The node ID of the node that originated this request. This is either the source node of the VP or a proxy node.
Target Node	Node ID of the target node of the restore path request. This is either the destination node of the VP or a proxy node.
Received From	The neighbor from which we received this message.
First Sequence Number	Sequence number of the first received copy of the corresponding restore-path request.
Last Sequence Number	Sequence number of the last received copy of the corresponding restore-path request.
Bandwidth	Requested bandwidth.
QoS	Requested QoS.
Timer	Used by the node to timeout the RPR.
T-Bit	Set to 1 when a Terminate indicator is received from any of the neighbors.
Pending Replies	Number of the neighbors that haven't acknowledged this message yet.
Sent To	<p>A list of all neighbors that received a copy of this message. Each entry contains the following information about the neighbor:</p> <p>AckReceived: Indicates if a response has been received from this neighbor.</p> <p>F-Bit: Set to 1 when Flush indicator from this neighbor.</p>

Table 1. RPR Fields

When an RPR packet arrives at a tandem node, a decision is made as to which neighbor should receive a copy of the packet. The choice of neighbors is related to variables such as link capacity and distance. Specifically, a particular neighbor is selected to receive a copy of the packet if:

1. The output link has enough resources to satisfy the requested bandwidth. ✓
2. The path through the neighbor is less than MAX_HOPS in length. In other words, the distance from this node to the target node is less than MAX_HOPS minus the distance from this node to the origin node.
3. The node hasn't returned a Flush response for this specific instance of the RPR, or a Terminate response for this or any other instance.

Fig. 5 illustrates the actions performed by tandem nodes in processing received RPR tests. Assuming that this is the first instance of the request, the node allocates the requested bandwidth on eligible links and transmits a modified copy of the received message onto them. The bandwidth remains allocated until a response (either positive or negative) is received from the neighboring node, or a positive response is received from any of the other neighbors (see Table 8 below). While awaiting a response from its neighbors, the node cannot use the allocated bandwidth to restore other VPs, regardless of their priority.

Processing of RPRs begins at step 500, in which the target node's ID is compared to the local node's ID. If the local node's ID is equal to the target node's ID, the local node is the target of the RPR and must process the RPR as such. This is illustrated in Fig. 5 as step 505 and is the subject of the flow diagram illustrated in Fig. 6. If the local node is not the target node, the RPR's HOP_COUNT is compared to MAX_HOP in order to determine if the HOP_COUNT has exceeded or will exceed the maximum number of hops allowable (step 510). If this is the case, a negative acknowledgment (NAK) with a Flush indicator is then sent back to the originating node (step 515). If the HOP_COUNT is still within acceptable limits, the node then determines whether this is the first instance of the RPR having been received (step 520). If this is the case, a Restore-Path Request Entry (RPRE) is created for the request (step 525). This is done by creating the RPRE and setting the RPRE's fields, including starting a time-to-live (TTL) or deletion timer, in the following manner:

```

RPRE.SourceNode = Header.Origin
RPRE.Destination Node = Header.Target
RPRE.FirstSequence Number = Header.SequenceNumber
RPRE.Last Sequence Number = Header.Sequence Number
RPRE.QoS = Header.Parms.RestorePath.QoS
RPRE.Bandwidth = Header. Parms.RestorePath.Bandwidth
RPRE.ReceivedFrom = Node ID of the neighbor that sent us this message
StartTimer (RPRE.Timer, RPR_TTL)

```


5

10

- 15

20

neighbors” refers to all adjacent nodes that are connected through links that meet the link-eligibility requirements previously described. Preferably, bandwidth is allocated only once for each request so that subsequent transmissions of the request do not consume any bandwidth.

- 5 Note that the bandwidth allocated for a given RPR is released differently depending on the type of response received by the node and the setting of the Flush and Terminate indicators in its header. Table 2 shows the action taken by a tandem node upon receiving a restore path response from a neighbor.

10

Response Type	Flush Indicator?	Terminate Indicator?	Received Sequence Number	Action
X	X	X	Not Valid	Ignore response
Negative	No	No	1 = Last	Ignore response
Negative	X	No	= Last	Release bandwidth allocated for the VP on the link on which the response was received.
Negative	Yes	No	Valid	Release bandwidth allocated for the VP on the link on which the response was received.
Negative	X	Yes	Valid	Release all bandwidth allocated for the VP.
Positive	X	X	Valid	Commit bandwidth allocated for the VP on the link the response was received on; release all other bandwidth.

Table 2. Actions taken by a tandem node upon receiving an RPR.

The Processing of Received RPRs by Destination or Target Node

- 15 Fig. 6 illustrates the process performed at the target node once the RPR finally reaches that node. When the RPR reaches its designated target node, the target node begins processing of the RPR by first determining whether this is the first instance of this RPR that has been received (step 600). If that is not the case, a NAK is sent with a Terminate indicator sent to the originating node (step 605). If this is the first
- 20 instance of the RPR received, the target node determines whether or not the VP

5 If the VP specified in the RPR terminates at this node (i.e. this node is indeed the target node), the target node determines whether an RPRE exists for the RPR received (step 615). If an RPRE already exists for this RPR, the existing RPRE is updated (e.g., the RPRE's LastSequenceNumber field is updated) (step 620) and the RPRE deletion timer is restarted (step 625). If no RPRE exists for this RPR in the target node (i.e., if this is the first copy of the instance received), an RPRE is created (step 630). Pertinent information from the RPR is copied into the RPRE (step 635), the bandwidth requested in the RPR is allocated on the input link by the target node (step 640) and an RPRE deletion timer is started (step 645). In either case, once the RPRE is either updated or created, a checksum is computed for the RPR (step 650) and written into the checksum field of the RPR (step 655). The RPR is then returned as a positive response to the origin node (step 660). The local (target) node then initiates its own matrix configuration. It will be noted that the RPRE created is not strictly necessary but helps to ensure that the processing of RPRs is consistent across nodes.

20

25 Negative RPR responses are processed as depicted in Fig. 7. An originating node begins processing a negative RPR response by determining whether the node has an RPRE associated with the RPR (step 700). If the receiving node does not have an RPRE for the received RPR response, the RPR response is ignored (step 705). If an associated RPRE is found, the receiving node determines whether the node sending
30 the RPR response is listed in the RPRE (e.g., is actually in the SentTo list of the RPRE) (step 710). If the sending node is not listed in the RPRE, again the RPR response is ignored (step 705).

5

20

25

15

20

25

30

1. Its HOP COUNT reaches the maximum allowed (i.e. MAX_HOPS).

- 5

10

15

25

30

Depending on the distance between the two nodes, the speed of link that connects them, and the latency of the equipment along the path, the timeout period can range

The above strategy of handling link errors can be improved upon, however, the fast restoration times required dictate that 2-way, end-to-end communication be carried out in less than 50ms. A limitation of the above-described solution is the time wasted while waiting for an acknowledgment to come back from the receiving node.

This problem is addressed in one embodiment by taking advantage of the multiple communication channels (i.e. OC-48's) that exist between nodes to:

- Preferably, the amount of packets sent is carefully controlled that the
20 broadcast does not create congestion in the network. The link efficiency is improved
further by using small packets during the restoration procedure. The present invention
can be practiced on a network node/element such as that described below.

- 23 -

Line™, Prodigy™, CompuServe™ and the like, by executing application specific software, commonly referred to as a browser, on one of client terminals 912(1)-(N).

One or more of client terminals 912(1)-(N) and/or one or more of servers 910(1)-(N) may be, for example, a computer system of any appropriate design, in general, including a mainframe, a mini-computer or a personal computer system. Such a computer system typically includes a system unit having a system processor and associated volatile and non-volatile memory, one or more display monitors and keyboards, one or more diskette drives, one or more fixed disk storage devices and one or more printers. These computer systems are typically information handling systems, which are designed to provide computing power to one or more users, either locally or remotely. Such a computer system may also include one or a plurality of I/O devices (i.e. peripheral devices) which are coupled to the system processor and which perform specialized functions. Examples of I/O devices include modems, sound and video devices and specialized communication devices. Mass storage devices such as hard disks, CD-ROM drives and magneto-optical drives may also be provided either as an integrated or peripheral device. One such exemplary computer system discussed in terms of client terminals 912(1)-(N) is shown in detail in Fig. 10.

Fig. 10 depicts a block diagram of a host computer system 1000 suitable for implementing the present invention, and exemplary of one or more of client terminals 912(1)-(N). Host computer system 1000 includes a bus 1012 which interconnects major subsystems of host computer system 1000. These subsystems include a central processor 1014, a system memory 1016 (typically RAM, but which may also include ROM, flash RAM, or the like), an input/output controller 1018, an external audio device such as a speaker system 1020 via an audio output interface 1022, an external device such as a display screen 1024 via display adapter 1026, serial ports 1028 and 1030, a keyboard 1032 (interfaced with a keyboard controller 1033), a storage interface 1034, a floppy disk drive 1036 operative to receive a floppy disk 1038, and a CD-ROM drive 1040 operative to receive a CD-ROM 1042. Also included are a mouse 1046 (or other point-and-click device, coupled to bus 1012 via serial port 1028), a modem 1047 (coupled to bus 1012 via serial port 1030) and a network interface 1048 (coupled directly to bus 1012).

Bus 1012 allows data communication between central processor 1014 and system memory 1016, which may include both read only memory (ROM) or flash memory (neither shown), and random access memory (RAM) (not shown), as previously noted. The RAM is generally the main memory into which the operating system and application programs are loaded and typically affords at least 16 megabytes of memory space. The ROM or flash memory may contain, among other code, the Basic Input-Output system (BIOS) which controls basic hardware operation such as the interaction with peripheral components. Applications resident with host computer system 1000 are generally stored on and accessed via a computer readable medium, such as a hard disk drive (e.g., fixed disk 1044), an optical drive (e.g., CD-ROM drive 1040), floppy disk unit 1036 or other storage medium. Additionally, applications may be in the form of electronic signals modulated in accordance with the application and data communication technology when accessed via network modem 1047 or interface 1048.

Storage interface 1034, as with the other storage interfaces of host computer system 1000, may connect to a standard computer readable medium for storage and/or retrieval of information, such as a fixed disk drive 1044. Fixed disk drive 1044 may be a part of host computer system 1000 or may be separate and accessed through other interface systems. Many other devices can be connected such as a mouse 1046 connected to bus 1012 via serial port 1028, a modem 1047 connected to bus 1012 via serial port 1030 and a network interface 1048 connected directly to bus 1012. Modem 1047 may provide a direct connection to a remote server via a telephone link or to the Internet via an internet service provider (ISP). Network interface 1048 may provide a direct connection to a remote server via a direct network link to the Internet via a POP (point of presence).

Many other devices or subsystems (not shown) may be connected in a similar manner. Conversely, it is not necessary for all of the devices shown in Fig. 10 to be present to practice the present invention. The devices and subsystems may be interconnected in different ways from that shown in Fig. 10. The operation of a computer system such as that shown in Fig. 10 is readily known in the art and is not discussed in detail in this application. Code to implement the present invention may

It will be noted that the variable identifier “N” is used in several instances in Fig. 10 to more simply designate the final element (e.g., servers 910(1)-(N) and client terminals 912(1)-(N)) of a series of related or similar elements (e.g., servers and client terminals). The repeated use of such variable identifiers is not meant to imply a correlation between the sizes of such series of elements, although such correlation may exist. The use of such variable identifiers does not require that each series of elements have the same number of elements as another series delimited by the same variable identifier. Rather, in each instance of use, the variable identified by “N” may hold the same or a different value than other instances of the same variable identifier.

The foregoing described embodiment wherein the different components are contained within different other components (e.g., the various elements shown as components of host computer system 1000). It is to be understood that such depicted architectures are merely exemplary, and that in fact many other architectures can be implemented which achieve the same functionality. In an abstract, but still definite

5

10

20